

IMAGE STABILIZATION SYSTEM AND METHOD FOR A VIDEO CAMERA

BACKGROUND OF THE INVENTION

1. Field of the Invention.

[0001] The present invention relates to a video camera stabilization system, and, more particularly, to a video camera stabilization system that counteracts unwanted and irregular image motion.

2. Description of the Related Art.

[0002] There are numerous known video surveillance systems which may be used to track a moving object such as a person or vehicle. Some such systems utilize a fixed camera having a stationary field of view (FOV). To fully cover a given surveillance site with a fixed camera system, however, it will oftentimes be necessary to use a significant number of fixed cameras.

[0003] Movable cameras which may pan, tilt and/or zoom may also be used to track objects. The use of a PTZ (pan, tilt, zoom) camera system will typically reduce the number of cameras required for a given surveillance site and also thereby reduce the number and cost of the video feeds and system integration hardware such as multiplexers and switchers associated therewith. Control signals for directing the pan, tilt, zoom movements typically originate from a human operator via a joystick or from an automated video tracking system.

[0004] A problem is that, in addition to being intentionally moved as a result of the PTZ control signals, the camera can also be unintentionally moved by external forces, such as air movement or mounting arm vibration. The inadvertent camera movements typically lead to a video image that is "jumpy", and thus unpleasant to observe.

[0005] What is needed in the art is a video system capable of stabilizing an image to compensate for inadvertent movement of the camera.

SUMMARY OF THE INVENTION

[0006] The present invention provides a video image stabilization system having a camera that captures an image that is larger than the associated image that is displayed on a screen. The system can compensate for inadvertent camera movement, i.e., remove the jumpiness from the displayed image, by adjusting the relative position of the portion of the image that is selected for display based upon the intended change in the field of view of the camera and an analysis of two of the captured images.

[0007] The invention comprises, in one form thereof, a video image stabilization system that includes a camera including an image capturing device configured to capture a video image wherein the camera has a selectively adjustable field of view. At least one processing device is operably coupled to the system wherein the processing device receives signals indicative of the field of view of the camera and images captured by the camera. The processing device sequentially grabs a first one of the images and a second one of the images and determines a stabilizing adjustment for the video image as a function of 1) an intended change in the field of view of the camera during an interval between the capture of the first image and the second image based upon the signals indicative of the field of view and 2) an analysis of the first and second images.

[0008] A display portion of each image captured by the camera may be selected for display wherein the selected display portion of each image is less than the entire captured image and the stabilizing adjustment determined by the processing device involves adjusting the relative location of the selected display portion within the captured image.

[0009] The invention comprises, in another form thereof, a video image stabilization system including a video camera having an image-capturing device configured to capture images. Each captured image is associated with a field of view and the camera has at least one selectively adjustable parameter wherein adjustment of the at least one camera parameter varies the field of view of the camera. The system also includes a display device configured to display a selected portion of images captured by the camera and a processing device operably coupled to the camera and to the display device. The processing device receives signals indicative of the at least one camera parameter and images captured by the camera. The processing device is operable to compensate for inadvertent movement of the video camera by adjusting the selected portion of the images displayed by the display device based upon the signals indicative of the at least one camera parameter and an analysis of a current captured image and a previously captured image.

[0010] The invention comprises, in another form thereof, a method of stabilizing a video image. The method includes capturing a plurality of images with a video camera, the video camera being mounted on a stationary support and having an adjustable field of view, and adjusting the field of view of the camera. The method also includes grabbing a first image and a subsequent second image captured by the video camera, displaying a selected display portion of each of the first and second images, the selected display portion of each image being less than the entire captured image and determining a stabilizing adjustment for the

second image. The stabilizing adjustment is determined as a function of an intended change in the field of view of the camera during an interval between the capture of the first image and the capture of the second image and an analysis of the first and second images wherein the stabilizing adjustment involves adjusting a relative location of the selected display portion within the second image.

[0011] The analysis of the first and second images may include transforming one of the first and second images wherein the scale of the one image is adjusted. In one embodiment, the analysis of the first and second images includes aligning one of the first and second images with the other of the first and second images based upon signals indicative of the field of view received from a system in operative communication with the camera and determining if the aligned images indicate the occurrence of unintentional camera movement. An image difference of the aligned images is used to determine if the aligned images indicate the occurrence of unintentional camera movement. A histogram of the image difference of the aligned images is generated and determining whether unintentional camera movement has occurred is accomplished by identifying a maximum peak in the histogram and a location of the maximum peak. When the maximum peak is not substantially centered on zero, unintentional motion is determined to have occurred and when such unintentional motion is detected, determination of the stabilizing adjustment includes identifying adjustments that minimize the image difference.

[0012] The invention comprises, in yet another form thereof, a method of stabilizing a video image. The method includes intentionally varying a field of view of a video camera by use of control signals, capturing a plurality of images with the camera, each of the images being associated with a field of view, displaying a selected portion of each of the plurality of images and adjusting the selected portion that is displayed for at least one of the plurality of images. The adjusting step compensates for inadvertent movement of the video camera and is based upon an intended field of view of the camera and an analysis of the at least one of the plurality of images and a second one of the images.

[0013] An advantage of the present invention is that it provides video images which are stabilized to limit or eliminate the distracting and unpleasant effects caused by inadvertent movements of the camera.

[0014] Another advantage of the present invention is that it allows for the stabilization of images acquired while the camera is in motion. Thus, the present invention can compensate for unintentional movement of a pan, tilt, zoom camera, such as that caused by support arm

vibration or the effect of wind on the camera housing, even while the camera is actively and purposefully changing its field of view.

BRIEF DESCRIPTION OF THE DRAWINGS

[0015] The above mentioned and other features and objects of this invention, and the manner of attaining them, will become more apparent and the invention itself will be better understood by reference to the following description of an embodiment of the invention taken in conjunction with the accompanying drawings, wherein:

Figure 1 is a schematic view of a video surveillance system in accordance with the present invention.

Figure 2 is a schematic view illustrating a change in a field of view of the camera of Figure 1 and the resulting images captured by the camera.

Figure 3 is a schematic view of a portion of the captured images of Figure 2.

Figure 4 is a block diagram of the processing device of Figure 1.

Figure 5 is a flowchart representing one embodiment of the operation of the video surveillance system of Figure 1.

Figure 6 is a histogram used in the analysis of two images.

[0016] Corresponding reference characters indicate corresponding parts throughout the several views. Although the exemplification set out herein illustrates an embodiment of the invention, in one form, the embodiment disclosed below is not intended to be exhaustive or to be construed as limiting the scope of the invention to the precise form disclosed.

DESCRIPTION OF THE PRESENT INVENTION

[0017] In accordance with the present invention, a video surveillance system 20 is shown in Figure 1. System 20 includes a camera 22 which is located within a partially spherical enclosure 24 and mounted on support 25. Stationary support 25 may take many forms, such as an outwardly extending support arm extending from an exterior edge of a building which may subject the supported camera to unintentional movement resulting from wind, vibrations generated by the camera motors, nearby machinery or a myriad of other sources. Enclosure 24 is tinted to allow the camera to acquire images of the environment outside of enclosure 24 and simultaneously prevent individuals in the environment being observed by camera 22 from determining the orientation of camera 22. Camera 22 includes a controller and motors which provide for the panning, tilting and adjustment of the focal length of camera 22. Panning movement of camera 22 is represented by arrow 26, tilting movement of camera 22 is represented by arrow 28 and the changing of the focal length of the lens 23 of camera 22,

i.e., zooming, is represented by arrow 30. As shown with reference to coordinate system 21, panning motion may track movement along the x axis, tilting motion may track movement along the y-axis and focal length adjustment may be used to track movement along the z-axis. In the illustrated embodiment, camera 22 and enclosure 24 are a Phillips AutoDome® Camera Systems brand camera system, such as the G3 Basic AutoDome ® camera and enclosure, which are available from Bosch Security Systems, Inc. formerly Phillips Communication, Security & Imaging, Inc. having a place of business in Lancaster, Pennsylvania. A camera suited for use with present invention is described by Sergeant et al. in U.S. Pat. No. 5,627,616 entitled Surveillance Camera System which is hereby incorporated herein by reference.

[0018] System 20 also includes a head end unit 32. Head end unit 32 may include a video switcher or a video multiplexer 33. For example, the head end unit may include an Allegiant brand video switcher available from Bosch Security Systems, Inc. formerly Phillips Communication, Security & Imaging, Inc. of Lancaster, Pennsylvania such as a LTC 8500 Series Allegiant Video Switcher which provides inputs for up to 64 cameras and may also be provided with eight independent keyboards and eight monitors. Head end unit 32 includes a keyboard 34 and joystick 36 for operator input. Head end unit 32 also includes a display device in the form of a monitor 38 for viewing by the operator. A 24 volt a/c power source 40 is provided to power both camera 22 and a processing device 50 that is operably coupled to both camera 22 and head end unit 32.

[0019] Illustrated system 20 is a single camera application, however, the present invention may be used within a larger surveillance system having additional cameras which may be either stationary or moveable cameras or some combination thereof to provide coverage of a larger or more complex surveillance area. One or more VCRs or other form of analog or digital recording device may also be connected to head end unit 32 to provide for the recording of the video images captured by camera 22 and other cameras in the system.

[0020] The general concepts underlying the operation of an image stabilization system in accordance with the present invention will now be discussed with reference to Figures 2 and 3. Camera 22 includes an image-capturing device such as a charge coupled device (CCD) that acquires a four-sided CCD video image 80 defined by CCD image boundary 82. Processing device 50 identifies or selects a central portion of CCD image 80 to be displayed on a screen of monitor 38 for viewing by an operator of system 20. More particularly, processing device 50 identifies a display image 84 defined by display image boundary 86 for

viewing on the screen of monitor 38. In one embodiment, the selected portion of the CCD video image, i.e., display image 84, includes approximately between 60% and 90% of the CCD video image. In a preferred embodiment, CCD image 80 can be approximately 30% larger than a display image 84 defined by display image boundary 86. In other words, in a preferred embodiment, approximately 23% of CCD image 80 is not displayed on the screen of monitor 38. However, for ease of illustration, CCD image 80 is shown in Figure 2 to be approximately between 100% and 200% larger than display image 84. By utilizing a display image that is smaller than the CCD image, processing device 50 can adjust the positions of the boundaries defining the display image relative to the entire CCD image for sequentially acquired images in order to counteract the effects of support arm vibration and other unintended movements of the camera, as described in more detail below.

[0021] After a period of time, the field of view (FOV) of camera 22 changes such that a second four-sided CCD image 88 is acquired. A second CCD image boundary 90 defines the limits of the CCD image 88. Figure 2 schematically illustrates the two images 80, 88 having different fields of view acquired by camera 22 and the extent to which the two images overlap.

[0022] The change in the FOV of camera 22 can be due to both intended and unintended, i.e., inadvertent, movements of camera 22. The intended movement of camera 22 that occurs can be due to panning, tilt and zoom control signals sent to camera 22 from head end unit 32. That is, the field of view of camera 22 can be varied by use of the control signals. Any unintended or inadvertent movement of camera 22 may be due to external forces, such as air movement or mounting arm vibration.

[0023] An intermediate, unadjusted display image 92 from CCD image 88, defined by an intermediate, unadjusted display image boundary 94, is identified by processing device 50 but is not displayed on the screen of monitor 38. Unadjusted display image 92 includes a portion of CCD image 88 that corresponds to the portion of CCD image 80 that is included in display image 84. That is, the relative position of display image 92 and CCD image boundary 90 is the same as the relative position of display image 84 to to CCD image boundary 82 with the distances between the display image boundaries and the corresponding CCD image boundaries being approximately equal for display images 84 and 92. For example, display images 84 and 92 may each be centered in their respective CCD images 80, 88.

[0024] Figures 2 and 3 schematically illustrate two images acquired by camera 22 and the effects and relationship between the intended motion of the camera, the actual motion of the

camera and the unintentional motion of the camera with regard to these two images. A common stationary feature 96 is identified in each of these images to clarify the relative movement of the FOV that has taken place between the acquisition of the two images. The same stationary feature 96 is shown in both solid and dashed lines in Figure 2 to indicate that feature 96 is present in both CCD image 80 and CCD image 88. The actual horizontal translation ΔP (generally corresponding to panning movement of camera 22) and vertical translation ΔT (generally corresponding to tilting movement of camera 22) of feature 96 relative to the CCD image boundaries 82, 90 that actually took place in the time period between the acquisition of the two CCD images 80, 88 is schematically represented in Figure 2.

[0025] In Figure 2, the positions CCD images 80 and 88 are depicted relative to stationary feature 96. As can be seen in Figure 2, images 80 and 88 overlap to a considerable extent but the FOV of camera 22 has changed between the acquisition of CCD image 80 to the more recently captured CCD image 88. The change in the FOV of camera 22 is best shown in Figure 3, wherein display images 84, 92 are superimposed on one another as they would successively appear on the screen of monitor 38 without any adjustment of display image boundaries 86, 94 for image stabilization purposes.

[0026] The vertical, i.e., in the tilt direction, change in position ΔT , and the horizontal, i.e., in the panning direction, change in position ΔP of feature 96 relative to the display image boundary 86, 94 is representative of the vertical (tilt) movement and horizontal (pan) movement of camera 22 that occurred between the capturing of CCD images 80 and 88 by camera 22 and is equivalent to the change in position of feature 96 relative to CCD image boundaries 82 and 90.

[0027] In other words, ΔT is equivalent to the difference between vertical distances v_0 and v_1 between feature 96 and corresponding horizontal edges of display image boundaries 86, 94, respectively. That is, $\Delta T = |v_0 - v_1|$. Similarly, ΔP is equivalent to the difference between horizontal distances h_0 and h_1 between feature 96 and corresponding vertical edges of display image boundaries 86, 94, respectively. That is, $\Delta P = |h_0 - h_1|$. Alternatively, ΔT may be said to be equivalent to the difference between vertical distances v_2 and v_3 between feature 96 and corresponding horizontal edges of CCD image boundaries 82, 90, respectively. That is, $\Delta T = |v_2 - v_3|$. Similarly, ΔP is equivalent to the difference between horizontal distances h_2 and h_3 between feature 96 and corresponding vertical edges of CCD image boundaries 82, 90, respectively. That is, $\Delta P = |h_2 - h_3|$.

[0028] The determination of the values of ΔT and ΔP is performed by processing device 50 based upon an analysis of the content of the two CCD images 80, 88. Matching techniques which may be used to register or align two images having overlapping content are well known in the art and used in a wide variety of computer vision and video processing applications. Two well-known methods are normalized correlation-based template matching and feature matching. For example, the identification of a stationary object such as feature 96 in each of the two images might be accomplished using a feature matching approach. One well known method of identifying features in video images is a corner detection method which analyzes the images to identify locations, or corners, in the image where there is a relatively sharp change in the intensity level of nearby pixels. These corners are then compared and matched in the two separate images.

[0029] Although such a feature matching approach can be employed with the present invention, the illustrated embodiment employs a template matching approach instead. Various template matching approaches are known in the art. Two well known approaches to template matching are a sum of square differences approach and a normalized cross-correlation approach. A sum of square differences approach typically relies upon the following mathematical relationship:

$$E = \sum_i \sum_j \|I(x'_i, y'_j) - M(x_i, y_j)\|^2, \quad (1)$$

where E is the total error, I is the image, M is the model template, (x, y) are the model image coordinates, and (x', y') are the transformed image coordinates. To register the two images, the total error E is minimized. In the example of Figure 2, first display image 84 corresponds to the model template, M , the difference between x and x' corresponds to ΔP and the difference between y and y' corresponds to ΔT . The relation between the transformed image coordinates and the model image coordinates is defined by a motion model. Different motion models can be used and an affine transformation as presented below models rotation, shearing, and translation:

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} a_0 & a_1 & a_2 \\ a_3 & a_4 & a_5 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \quad (2)$$

This motion model (2) requires at least 3 point pairs between the current image (second CCD image 88) and the model image (first display image 84).

[0030] Although a sum of square differences approach can be used with the present invention, the illustrated embodiment utilizes a normalized cross-correlation approach instead. Template matching using cross-correlation is based upon squared Euclidean distances:

$$d_{f,t}^2(u,v) = \sum_{x,y} [f(x,y) - t(x-u, y-v)]^2 \quad (3)$$

Wherein f is the image, e.g., second CCD image 88, and the summation is over x, y under the window, e.g., first display image 84, containing a feature t positioned at (u, v) . The expansion of d^2 provides:

$$d_{f,t}^2(u,v) = \sum_{x,y} [f^2(x,y) - 2f(x,y)t(x-u, y-v) + t^2(x-u, y-v)] \quad (4)$$

The term $\sum t^2(x-u, y-v)$ is constant. If the term $\sum f^2(x,y)$ is approximately constant then a measure of the similarity between the image and the feature is provide by the remaining cross-correlation term:

$$c(u,v) = \sum_{x,y} f(x,y)t(x-u, y-v) \quad (5)$$

There are several disadvantages, however, to using this term for template matching and, oftentimes, a correlation coefficient is used instead. The correlation coefficient may be obtained by normalizing the image and feature vectors to unit length, providing:

$$\gamma(u,v) = \frac{\sum_{x,y} [f(x,y) - \bar{f}_{u,v}] [t(x-u, y-v) - \bar{t}]}{\left\{ \sum_{x,y} [f(x,y) - \bar{f}_{u,v}]^2 \sum_{x,y} [t(x-u, y-v) - \bar{t}]^2 \right\}^{0.5}} \quad (6)$$

where \bar{t} is the mean of the feature and $\bar{f}_{u,v}$ is the mean of $f(x,y)$ in the region under the feature. This relationship, known to those having ordinary skill in the art, is commonly referred to as normalized cross-correlation. By maximizing the normalized correlation between the first display image 84 and the second CCD image 88, processing device 50 can

determine a value for ΔT and ΔP .

[0031] In addition to determining the actual change in the field of view as indicated by ΔT and ΔP , processing device 50 also determines the intended change, if any, in the field of view, FOV, of camera 22. In Figure 2, the intended change in the FOV in the vertical (tilting) and horizontal (panning) directions respectively correspond to ΔT_I and ΔP_I . For example, when the camera is intentionally panned or tilted, there will be a change of the FOV of the camera. In Figure 2, this intentional change in the FOV of the camera corresponds to ΔT_I and ΔP_I wherein ΔT_I is a vertical image distance corresponding to the intentional tilting of camera 22 and ΔP_I is a horizontal image distance corresponding to the intentional panning of camera 22.

[0032] In the illustrated embodiment, camera 22 communicates to processing device 50 frame-based synchronized information on the pan, tilt and zoom position of camera 22 with each image acquired by camera 22. Processing device 50 utilizes the pan, tilt and zoom settings associated with each of the two images 80, 88 to determine the intended change, if any, in the pan, tilt and zoom positions of camera 22 and, thus, the magnitude and direction of ΔT_I and ΔP_I .

[0033] The intended change in the FOV can also be determined by alternative methods. For example, processing device 50 can analyze the panning, tilt and zoom control signals sent to camera 22 to determine the intended change in the FOV. Such control signals may originate from head end unit 32, e.g., due to the manual manipulation of joystick 36, and be communicated to processing device 50 directly or through camera 22. Alternatively, the camera control signals may originate from processing device 50, e.g., camera control signals generated by an automatic intruder tracking program. Alternatively, the camera control signals may be originated by another source and communicated to processing device 50 or processing device 50 may determine the intended change in the display image by analyzing other data gathered from system 20.

[0034] Once both the actual change in the FOV of the camera, as represented by ΔT and ΔP , and the intended change in the FOV of the camera, as represented by ΔT_I and ΔP_I , are known, it can be determined whether there has been any unintentional movement of the camera in either or both of the vertical (tilting) and horizontal (panning) directions. The values ΔT_E and ΔP_E represent the difference between the actual change in the FOV (ΔT and ΔP) and the intended change in the FOV (ΔT_I and ΔP_I) for the two sequential images, i.e., the change in the FOV that is due to the inadvertent movement of camera 22. Processing device

50 compensates for such inadvertent movement of camera 22 by adjusting the relative position of the display image within the CCD image based upon the values of ΔT_E and ΔP_E . In Figure 2, an adjusted display image 98 defined by adjusted display boundary 100 would be displayed on the screen of monitor 38.

[0035] When communicating images to head end unit 32 for display on monitor screen 38, processing device 50 can clip the images and forward only the selected display image portion, e.g., display images 84, 98, or it may forward the entire CCD image, e.g., CCD images 80, 88, together with the coordinates of the display image boundaries, e.g., display boundaries 86, 100.

[0036] As long as the entire adjusted display image 98 is available on CCD image 88, adjusted display image 98 can be displayed on the screen of monitor 38. The outer edges of the selected portion of the video image, i.e., the display image boundary, and the outer edges of the CCD video image are separated by margins on all four sides. The adjustment of the display image is, thus, limited to width of the corresponding margin. That is, the extent to which the display image can be repositioned within the CCD image is limited by a threshold value corresponding to the size difference between the display image and the CCD image. The size of the margins may be selected based upon the anticipated magnitude of the vibrations or other unintended movements that will be encountered by the surveillance camera.

[0037] A key feature provided by the present invention is that it not only stabilizes the display image while the camera is stationary, but also provides for the stabilization of the display image while the FOV of the camera is being adjusted. As described above, in addition to the adjustable panning and tilt angles of camera 22, camera 22 also has an adjustable zoom setting. The preceding discussion concerning the determination of ΔT and ΔP ; ΔT_I and ΔP_I ; and ΔT_E and ΔP_E did not specifically address the situation where the change in the FOV involves a change in the zoom setting of camera 22, i.e., a change in the focal length of the lens of camera 22, in addition to panning or tilting movement of the camera. The vibration of camera 22 or of the structure supporting camera 22 is not likely to result in an unintentional change in the zoom setting of camera 22. As a result, the zoom setting communicated from camera 22 to processing device 50 is assumed to be correct.

[0038] An intentional and precise change in the zoom setting between the acquisition of images 80, 88, however, will impact the image stabilization process by altering the relative size of the content of images 80, 88. Alternative embodiments of the present invention may

address such a change in the zoom setting of camera 22 in alternative manners. The simplest approach is to only perform the image stabilization process when the two sequential images being analyzed by processing device 50 have the same zoom setting. This approach will suspend the operation of the image stabilization process when the focal length of camera 22 is being actively changed. For some applications where the focal length of the camera is subject to relatively few changes or where the camera does not include an adjustable focal length, this approach may not appreciably degrade the value of the image stabilization system. Moreover, when the focal length is being actively changed, the image is undergoing a noticeable transformation which is typically brief and wherein small unintentional image translations will not be as distracting as when the camera is stationary or being subjected to only panning and/or tilting movements.

[0039] It is also possible to continue to perform the image stabilization process as the focal length of camera 22 is being changed. If this approach is taken, at least one of the two images being compared, i.e., first and second CCD images 80, 88, must be transformed so that the objects depicted in each of the images are depicted at a common scale and can be aligned. To accomplish this, geometric transforms may be used to modify the position of each pixel within the image. Another way to think of this is as the moving of all pixels from one location to a new location based upon the camera motion. One such method for transforming a first image to align it with a second image wherein the camera was adjusted between the acquisition of the two images is discussed by Trajkovic in U.S. Pat. App. Pub. No. 2002/0167537 A1 entitled Motion-Based Tracking With Pan-Tilt-Zoom Camera which is hereby incorporated herein by reference.

[0040] Alignment of consecutive images acquired at different focal lengths requires translation as well as scaling and rotation of one image to align it with the previous image. Of these three operations translation is the simplest. Warping, a process in which each pixel is subjected to a general user-specified transformation, may be necessary to reduce, expand, or modify an image to a standard size before further processing can be performed. Images produced by such geometric operations are approximations of the original. The mapping between the two images, the current image, e.g., I_1 , and a reference image, e.g., I_2 , is defined by:

$$p' = sQRQ^{-1}p = Mp \quad (7)$$

where p and p' denote the homographic image coordinates of the same world point in the first

and second images, s denotes the scale image (which corresponds to the focal length of the camera), Q is the internal camera calibration matrix, and R is the rotation matrix between the two camera locations.

[0041] Alternatively, the relationship between the image projection coordinates p and p' , i.e., pixel locations (x,y) and (x', y') , of a stationary world point in two consecutive images may be written as:

$$x' = \frac{m_{11}x + m_{12}y + m_{13}}{m_{31}x + m_{32}y + m_{33}} \quad (8)$$

$$y' = \frac{m_{21}x + m_{22}y + m_{23}}{m_{31}x + m_{32}y + m_{33}} \quad (9)$$

Where $\begin{bmatrix} m_{ij} \end{bmatrix}_{3 \times 3}$ is the homography matrix M that maps (aligns) the first image to the second image.

[0042] The main task in such image alignment is to determine the matrix M . From equation (7), it is clear that given s , Q and R it is theoretically straightforward to determine matrix M . In practice, however, the exact values of s , Q , and R are often not known. Equation (7) also assumes that the camera center and the center of rotation are identical, which is typically only approximately true, however, this assumption may be sufficiently accurate for purposes of providing image stabilization. In the illustrated embodiment, camera 22 provides data, i.e., pan and tilt values for determining R and zoom values for determining s , on an image synchronized basis and with each image it communicates to processing device 50.

[0043] With this image specific data, the translation, rotation, and scaling of one image to align it with the second image can then be performed using the homographic method outlined above. In this method, a translation is a pixel motion in the x or y direction by some number of pixels. Positive translations are in the direction of increasing row or column index; negative ones are the opposite. A translation in the positive direction adds rows or columns to the top or left to the image until the required increase has been achieved. Image rotation is performed relative to an origin, defined to be at the center of the motion and specified as an angle. Scaling an image means making it bigger or smaller by a specified factor. The following approximations may be used to represent such translation, rotation and scaling:

$$\begin{aligned}x' &= s(x \cos \alpha - y \sin \alpha) + t_x \\y' &= s(y \sin \alpha + x \cos \alpha) + t_y\end{aligned}\tag{10}$$

wherein

s is the scaling (zooming) factor.

α is the angle of rotation about the origin;

t_x is the translation in the x direction; and

t_y is the translation in the y direction.

By introducing new independent variables $a_1 = s \cos \alpha$ and $a_2 = s \sin \alpha$, equation (10) becomes:

$$\begin{aligned}x' &= a_1 x - a_2 y + t_x \\y' &= a_2 x + a_1 y + t_y\end{aligned}\tag{11}$$

After determining a_1 , a_2 , t_x and t_y , the two images, I_1 and I_2 , can be aligned. If these values are obtained using zoom, pan and tilt values obtained directly from the camera, t_x and t_y will correspond to ΔP_1 and ΔT_1 . If the transformed images display global image motion and are not properly aligned by this process, unintentional movement has occurred and the values of ΔP_E and ΔT_E can be determined for these transformed images using normalized cross-correlation.

[0044] The hardware architecture of processing device 50 is schematically represented in Figure 4. In the illustrated embodiment, processing device 50 includes a system controller board 64 in communication with a power supply / IO board 66. A power line 42 connects power source 40 to converter 52 in order to provide power to processing device 50. Processing device 50 receives a raw analog video feed from camera 22 via video line 44, and video line 45 is used to communicate video images to head end unit 32. In the illustrated embodiment, video lines 44, 45 are coaxial, 75 ohm, 1 Vp-p and include BNC connectors for engagement with processing device 50. The video images provided by camera 22 can be analog and may conform to either NTSC or PAL standards. When processing device 50 is inactive, i.e., turned off, video images from camera 22 can pass through processing device 50 to head end unit 32 through analog video line 54, analog circuitry 68, analog video line 70 and communications plug-in board 72. Board 72 can be a standard communications board

capable of handling biphasic signals and including a coaxial message integrated circuit (COMIC) for allowing two-way communication over video links.

[0045] Via another analog video line 56, an analog to digital converter 58 receives video images from camera 22 and converts the analog video signal to a digital video signal. After the digital video signal is stored in a buffer in the form of SDRAM 60, the digitized video images are passed to video content analysis digital signal processor (VCA DSP) 62. The video stabilization algorithm described above with reference to Figures 2 and 3 is performed in VCA DSP 62. The adjusted display image is sent to digital to analog converter 74 where the video signal is converted to an analog signal. The resulting annotated analog video signal is sent via analog video lines 76 and 70 to communications plug-in board 72, which then sends the signal to head end unit 32 via video line 45.

[0046] In the illustrated embodiment, video input to system controller board 64 is limited to 1 Vp-p and if the video signal exceeds 1 Vp-p it will be clipped to 1 Vp-p. However, alternative embodiments having a greater or lesser capacity may also be employed with the present invention. Processor 62 may be a TriMedia TM-1300 programmable media processor available from Phillips Electronics North America Corporation. At start up, processor 62 loads a bootloader program. The boot program then copies the VCA application code from a memory device such as flash memory 78 to SDRAM 60 for execution. In the illustrated embodiment, flash memory 78 provides 1 megabyte of memory and SDRAM 60 provides eight megabytes of memory. Since the application code from flash memory 78 is loaded on SDRAM 60 upon start up, SDRAM 60 is left with approximately seven megabytes of memory for video frame storage.

[0047] In the embodiment shown in Figure 4, system controller board 64 is connected to communications plug-in board 72 via a biphasic digital data bus 102, an I2C data bus 104, and an RS-232 data bus 106. System controller board 64 is connected to an RS-232 / RS-485 compatible transceiver 108 via RS-232 data bus 110. A line 49, which can be in the form of an RS-232 debug data bus, communicates signals from head end unit 32 to processing device 50. The signals on line 49 can include signals that can be modified by processing device 50 before being sent to camera 22. Such signals may be sent to camera 22 via line 48 in communication with microprocessor 112. Microprocessor 112 can operate system controller software and may also be in communication with VCA components 114. Thus, VCA components such as VCA DSP 62 can send signals to camera 22 via microprocessor 112 and line 48.

[0048] System controller board 64 may also include a field programmable gate array 116 including a mask memory 118, a character memory 120, and an on screen display (OSD) memory 122. Similarly, VCA components 114 may include a mask memory 124, a character memory 126, and an on screen display (OSD) memory 128. These components may be used to mask various portions of the image displayed on screen 38 or to generate textual displays for screen 38. Finally, system controller board 64 can include a parallel data flash memory 130 for storage of user settings.

[0049] In the illustrated embodiment, the only necessary commands conveyed to processing device 50 that are input by a human operator are on/off commands, however, even these on/off commands may be automated in alternative embodiments. Such on/off commands and other serial communications are conveyed via bi-phase line 46 between head end unit 32 and camera 22, and between processing device 50 and camera 22 via line 48. In the illustrated embodiment, processing device 50 is provided with a sheet metal housing and mounted proximate camera 22, however, processing device 50 may also be mounted employing alternative methods and at alternative locations. Alternative hardware architecture may also be employed with processing device 50. Such hardware should be capable of running the software and processing at least approximately five frames per second for good results. It is also noted that by providing processing device 50 with a sheet metal housing its mounting on or near a PTZ camera is facilitated and system 20 may thereby provide a stand alone embedded platform which does not require a personal computer-based image stabilization system. If desired, however, the present invention may also be employed using a personal computer based system.

[0050] Processing device 50 can perform several functions, including capturing video frames acquired by camera 22, identifying a stationary feature in the video frames, determining the intended change in the camera FOV based upon signals sent to or received from camera 22, identifying a stationary feature and determining the actual change in the camera FOV, comparing the intended and actual change in the camera FOV to determine the magnitude of the image translations resulting from the unintentional motion of the camera and selecting display image coordinates to counteract the translations resulting from the unintentional motion of the camera. Processing device 50 may also be used to perform an automated tracking function. For example, processing device 50 may also provide an automated tracking system wherein processing device 50 is used to identify moving target objects in the FOV of the camera and then generate control signals which adjust the pan, tilt

and zoom settings of the camera to track the target object and maintain the target object within the FOV of the camera. As the pan, tilt and zoom settings of the camera are automatically adjusted to track the target object the images displayed by the system may be stabilized by utilizing an image stabilization system in accordance with the present invention. An example of an automated tracking system that may be employed by system 20 is described by Sablak et al. in U.S. Patent Application Serial No. 10/306,509 filed on November 27, 2002 entitled "VIDEO TRACKING SYSTEM AND METHOD" the disclosure of which is hereby incorporated herein by reference. An image stabilization system in accordance with the present invention may also be employed to stabilize images wherein the camera is being manually adjusted, such as by the manual manipulation of joystick 36 or in other applications.

[0051] Figure 5 provides a flow chart which graphically illustrates the general logic of an embodiment of the video stabilization algorithm used by system 20 as described above. As shown in Figure 5, after turning processing device 50 on, it is initialized at step 180 by copying the application code from flash memory 78 to SDRAM 60 for execution. Block 182 represents the remaining memory of SDRAM 60 which is available as a ring buffer for storage of video image frames and data associated with each frame received from camera 22 and stored for processing by processor 62. At decision block 184, processor 62 determines if the first flag is true. The first flag is true only when no images from camera 22 have been loaded to SDRAM 60 for analysis by processor 62. Thus, when processing device 50 is turned on, the first time decision block 184 is encountered, the first flag will be true and processor 62 will proceed to block 186. Block 186 represents the grabbing of two images by processor 62. Processor 62 then proceeds to block 188 where a smoothing filter is applied to each image. Application of the smoothing filter involves taking two sub-samples of the images. The first sub-sampling step creates a QCIF resolution sub-sample (i.e., an image having a quarter of the resolution of the original NTSC or PAL CIF resolution image) of the current I_1 and I_2 images. The sub-sampling process groups adjacent pixels together to define an average value for the grouped pixels. The purpose of the sub-sampling process is to reduce the time consumed by the computational processes involve in analyzing the image. A second sub-sample of the first sub-sample is then taken resulting in images having 1/16 the resolution of the original CIF resolution images to thereby further increase the speed of the image stabilization process. Block 190 represents the acquisition of the pan, tilt and focal length (i.e., zoom) values for each of the images from the camera. (This data represents the

intended values and does include the unintentional motion components of the pan and tilt values.) Although block 190 is shown at a location following the acquisition of the images, the pan, tilt and zoom data may be communicated by camera 22 to processing device 50 simultaneously with the images.

[0052] Next, block 192 represents the calculating of a homography matrix for rotation and scaling using the pan, tilt and zoom values of the current image. Then, in block 194, an image transformation is performed by homography to align the second image on to the first image. (If no pan, tilt or zoom adjustments to the camera have been made during the time interval between the acquisition of the two images, no transformation of the images is required.)

[0053] At block 196, the image difference of the two aligned images is calculated to determine if any motion is present in the images. Thus, after aligning images I_1 and I_2 , the image difference, Δ , is calculated in accordance with the following equation:

$$\Delta = | I_2 - I_1 | \quad (12)$$

A histogram of these differences is then calculated. The stationary background of the image will typically be responsible for generating the largest peak in the histogram. Thus, if the largest peak is centered around zero, the background of the two images is aligned and there has not been any unintentional camera movement during the interval between the acquisition of the two images. If this is the case, the process returns to decision box 184. If the largest peak in the histogram is not centered around zero, this indicates global motion, i.e., the background is in motion. This is interpreted as indicating unintentional motion of the camera and the process proceeds to box 198. It is also possible that a moving object will be located in the acquired images. If such a moving object is present in the images, the histogram will typically have two peaks associated with it, one corresponding to the background and one corresponding to the moving object. If this is the case, the background of the image is assumed to take up more area of the image than the moving object and the largest peak value is used to make the above determination. In other words, the histogram is evaluated to determine whether the largest peak of the histogram is centered on zero, indicating the absence of unintentional camera movement, or whether the largest peak indicates global movement, signifying the presence of unintentional camera movement. For example, Figure 6 schematically depicts a histogram of image differences wherein the primary peak is centered on zero, indicating no unintentional camera movement, and also including a

secondary peak 212 that is not centered on zero thereby indicating the presence of a moving object within the images.

[0054] If unintentional camera movement is detected and the system proceeds to block 198, the extent of the motion of the between the two images is determined using normalized cross-correlation (NCC). Because the transformation of the two images has already aligned the images in the manner predicted by the values representing ΔT_1 and ΔP_1 , the process step of determining the motion that has occurred between the two images at box 198 corresponds to determining the value of ΔT_E and ΔP_E .

[0055] Next, at block 200, the data for images I_1 and I_2 are swapped. The swapping of image data is done so that when a new image is grabbed and placed in the buffer after completing the translation of block 202, the new image and data associated therewith will overwrite the image and data associated with the older of the two images already present in the buffer.

[0056] Block 202 represents the translation of the display image boundaries on the CCD image of the most recently acquired image by amounts equivalent to ΔT_E and ΔP_E to thereby provide a stabilized video image.

[0057] The process then returns to block 184 where the first flag will no longer be true and the process will proceed to block 208 where a single new image will be grabbed and overwrite image I_2 in the buffer. Processor 62 then proceeds to block 210 where the smoothing filter is applied to the new image. Block 212 represents the acquisition of the pan, tilt and zoom settings of the camera at the time the new image was acquired by the camera. The process then proceeds to block 192.

[0058] In the exemplary embodiment, camera 22 is continually acquiring new images and the computational analysis performed by processing device 50 to compare the current image with a previously acquired image takes longer than the time interval between the individual images acquired by camera 22. When processing device 50 completes its analysis of a set of images, it grabs the most recently acquired image for the next analysis. Thus, there may be one or more images that are captured and communicated by camera 22 that were captured between two images that are sequentially grabbed by processing device 50 for analysis. As a result, when it is determined at block 198 (in the flow chart of Figure 5) that the display boundaries for the current image require adjustment to correct for the unintentional movement of the camera, several additional images will be acquired by camera 22 and displayed on monitor screen 38 before the next image stabilization analysis is completed.

Setting the location of the display boundary in these additional unanalyzed images can be done in different manners. For example, the display boundary used with these additional images utilize the adjusted boundary position, e.g., display boundary 100, that resulted from the analysis of the last two grabbed images could be employed with the unanalyzed images until the stabilization process positively determined that the display boundary should be readjusted. Alternatively, after altering the display boundary for the one image which was specifically analyzed and determined to have been influenced by unintentional camera movement, the display boundary could be immediately returned to its centered location in the CCD display. In other words, the display boundary would remain at its centered location for all images except for those specific images which have been grabbed, analyzed and determined to have been affected by unintentional camera movement and for which an altered display image boundary is specifically determined. It is also possible for the display boundary to be returned to the centered location in the CCD image boundary in incremental amounts, over several sequentially acquired images, after the display boundary has been adjusted to account for unintentional camera movement.

[0059] In most surveillance camera applications, it is anticipated that the types of vibrations and other external forces acting on the camera will cause unintentional camera movements that have a very brief deviation from the intended position with possible oscillations about an equilibrium position corresponding to the intended position of the camera. In this type of situation, having the display boundary return immediately to the centered location can be beneficial. Returning to the centered position may not only may reflect the underlying unintentional camera motion and thus facilitate the display of a stabilized image, it also prevents migration of the display boundaries to the edge of the CCD display image which could limit the effectiveness of the image stabilization process and could result from a process which does not return the display image boundaries to a centered location in the CCD image.

[0060] While this invention has been described as having an exemplary design, the present invention may be further modified within the spirit and scope of this disclosure. This application is therefore intended to cover any variations, uses, or adaptations of the invention using its general principles.